Statistical mechanics analysis of the equilibria of linear economies

A De Martino†, M Marsili‡, I Pérez Castillo¶

† INFM–SMC and Dipartimento di Fisica, Università di Roma "La Sapienza", Piazzale Aldo Moro 2, 00185 Roma (Italy)
‡ The Abdus Salam ICTP, Strada Costiera 11, 34014 Trieste (Italy)
¶ Instituut voor Theoretische Fysica, Katholieke Universiteit Leuven, Celestijnenlaan 200D, 3001 Leuven (Belgium)

E-mail: andrea.demartino@roma1.infn.it, marsili@ictp.trieste.it, Isaac.Perez@fys.kuleuven.ac.be

Abstract. The optimal ('equilibrium') macroscopic properties of an economy with N industries endowed with different technologies, P commodities and one consumer are derived in the limit $N \to \infty$ with $n = N/P$ fixed using the replica method. When technologies are strictly inefficient, a phase transition occurs upon increasing n. For low n, the system is in an expanding phase characterized by the existence of many profitable opportunities for new technologies. For high n, technologies roughly saturate the possible productions and the economy becomes strongly selective with respect to innovations. The phase transition and other significant features of the model are discussed in detail. When the inefficiency assumption is relaxed, the economy becomes unstable at high n.

1. Introduction

In recent years it has become increasingly clear that the emergence of non-trivial collective phenomena in systems of heterogeneous interacting agents (e.g. markets, information networks, or ecosystems) presents deep analogies with the occurrence of phase transitions in disordered statistical systems, when heterogeneities in the former can be mapped onto disorder variables in the latter. The ideas and techniques of mean-field spin-glass theory [1] play a major role in the current studies of multi-agent systems. As happened in other interdisciplinary applications of spin-glass theory from neural networks [2] to computer science [3], the effectiveness of these methods in economics-inspired problems can go well beyond the purely technical level. Some issues like the emergence of collective phenomena in evolving populations of competing agents have already started to be elucidated via spin-glass type of ideas (see e.g. the literature on the Minority Game [4]). Another intriguing class of problems arising in economics where such a potential can be fully expressed concerns the characterization of the equilibria of heterogeneous economies.

In a nutshell, economic equilibria are optimal microscopic states of an economy in which (a) all economic agents maximize their respective payoffs (firms maximize profits and consumers maximize utilities) and (b) the prices of goods are such that the total demand and the total supply of each commodity match exactly [5, 6]. Microscopic degrees of freedom are typically the scale of production of firms and the level of

consumption of consumers, while technologies and endowments or utility functions are heterogeneous across firms and consumers respectively. The central problem is extracting robust macroeconomic properties and laws that can be tested against observed patterns of real economies. Heterogeneities complicate things considerably from a theoretical viewpoint. The traditional approach of economics is based on the so-called representative agent description, in which each class of economic agents (firms and consumers) is substituted by a single, "average" agent whose characteristics are, to a large extent, postulated. An alternative pathway consists in modeling heterogeneities through randomness [7]. While this approach has drawn heavily from tools of statistical mechanics, mainly via analogies with random field models [8], the possibility of systematically employing statistical mechanics to clarify the statistical structure of economic equilibria has only been partially explored thus far.

In this work, we apply spin-glass techniques to the study of a random production economy assuming that technologies (i) are random vectors, (ii) are to some degree inefficient and (iii) can be operated at any (non-negative) scale. Within this framework, which is standard in economics [9], it is possible to *derive* an effective-agent problem whose solution embodies that of the original multi-agent problem. Key macroeconomic quantities such as averages and distributions of prices, operation scales and consumptions can be computed exactly. Remarkably, they turn out to be either directly connected or easily derived from the spin-glass order parameters that naturally arise in calculations. Consequently, a macroeconomic picture emerges naturally from microeconomic interactions and the corresponding phase structure can be fully characterized.

Our emphasis in this paper will be mainly on 'physical' properties of the model (e.g. phase transitions), so we will not analyze the behavior of such quantities as the gross domestic product of the economy. For a more articulated economic discussion, we refer the reader to [10], where a slightly simpler model than the one presented here is studied in a fully economic perspective. The outline of this work is hence as follows. In Sec. 2 we expound the formal aspects of the multi-agent problem and outline the spin-glass approach. Two types of economies are considered: those with strictly inefficient technologies and those whose technologies are inefficient *on average*. These two situations give rise to strikingly different macroscopic properties. The former choice yields a more realistic economic picture and a non-trivial physical scenario, which is thoroughly analyzed in Sec. 3. In the latter case, the economy displays simpler but less realistic properties, that are the subject of Sec. 4. Finally, in Sec. 5, we fomulate our conclusions and discuss some possible directions for further work.

2. The problem and its solution

2.1. Statement of the problem

The definitions of our model parallel those employed in the theory of economic equilibria, see e.g. [9]. We consider an economy with N firms or industries and P commodities. Each industry is endowed with a technology or activity that allows the transformation of some commodities ("inputs") into others ("outputs"). Technologies are denoted as P-dimensional vectors $\boldsymbol{\xi}_i = \{\xi_i^\mu\} \in \mathbb{R}^P$, where negative (resp. positive) components represent quantities of inputs (resp. outputs). Furthermore, each firm can operate its technology at a scale $s_i \geq 0$, meaning that when run at scale s_i the

industry i produces or consumes a quantity $\xi_i^\mu s_i$ of commodity μ. The profit of firm i is $\pi_i = s_i(\boldsymbol{p}\cdot\boldsymbol{\xi}_i)$, where $\boldsymbol{p} = \{p^\mu\}$ is the price vector, which we assume to be non-negative. Each firm chooses s_i by solving

$$\max_{s_i \geq 0} \ \pi_i \tag{1}$$

at fixed prices. Produced goods are bought by consumers to satisfy their needs. Here we consider the case of a single consumer ('society') endowed with an initial bundle of goods $\boldsymbol{y} = \{y^\mu\}$, $y^\mu \geq 0$, and with a prescribed utility function $U(\boldsymbol{x})$ ('welfare'). The consumer selects the desired consumption $\boldsymbol{x} = \{x^\mu\}$ by solving

$$\max_{\boldsymbol{x}\in B} \ U(\boldsymbol{x}), \qquad B = \{\boldsymbol{x} \geq 0 : \boldsymbol{p}\cdot\boldsymbol{x} \leq \boldsymbol{p}\cdot\boldsymbol{y}\} \tag{2}$$

at fixed prices. B denotes the consumer's budget set. We specialize here to the case of a separable utility:

$$U(\boldsymbol{x}) = \sum_{\mu=1}^{P} u(x^\mu) \tag{3}$$

which simplifies the analysis considerably. We assume in this way that commodities are *a priori* equivalent for the consumer. While our calculations hold in general for u's such that $u'(x) > 0$ and $u''(x) < 0$, for numerical calculations we set $u(x) = \log x$.

Equilibria are defined as microscopic configurations $(\boldsymbol{s}, \boldsymbol{x})$, with $\boldsymbol{s} = \{s_i\}$ and $\boldsymbol{x} = \{x^\mu\}$, that solve (1) and (2) and such that

$$\boldsymbol{x} = \boldsymbol{y} + \sum_{i=1}^{N} s_i \boldsymbol{\xi}_i \tag{4}$$

The above "market-clearing" condition ensures that the total demand of each good, $x^\mu - y^\mu$, matches the total supply $\sum_i \xi_i^\mu s_i$.

We concentrate on the limit of large economies, namely $N \to \infty$, $P \to \infty$, $n = N/P$ fixed. The consumer's initial endowments y^μ are assumed to be sampled independently from a distribution $\rho(y)$. The complexity of the production processes is modeled by postulating that technologies $\boldsymbol{\xi}_i$ are quenched random vectors. In order to avoid unrealistic features, it is crucial to impose that it is not possible to produce, by some combination of the available technologies, a positive amount of some commodity without consuming a positive amount of some other commodity ("impossibility of the Land of Cockaigne" [9]). We enforce this condition in two ways.

Hard-constrained technologies. The ξ_i^μ's are taken to be Gaussian random variables with zero mean and variance Δ_i/P satisfying

$$\forall i : \qquad \sum_{\mu=1}^{P} \xi_i^\mu = -\epsilon_i \tag{5}$$

The parameters $\epsilon_i > 0$ measure the inefficiency of the transformation process of technologies, i.e. the difference between the total quantity of inputs and the total quantity of outputs.

Soft-constrained technologies. The ξ_i^μ's are taken to be Gaussian random variables with mean $-\epsilon_i/P$ and variance Δ_i/P.

These two choices give rise to very different macroscopic properties. In both cases, we further assume that

$$\epsilon_i = \eta\sqrt{\Delta_i} \tag{6}$$

and that Δ_i is drawn from a distribution $g(\Delta)$ independently for each i.

The parameters that ultimately regulate the economy's equilibria are thus n and η, together with the functions $u(x)$, $\rho(y)$ and $g(\Delta)$.

As said above, in this setup commodities are statistically equivalent in all respects. Hence, this model focuses on the ability of the economy to produce scarce goods using abundant goods as inputs. Our aim is thus to understand how the global level of prosperity, measured by the utility of the consumer and by the amount of economic activity, increases as the repertoire of available technologies expands.

2.2. Statistical mechanics approach

In the following, we shall generally denote by brackets $\langle\cdots\rangle_z$ the average over a random variable z, irrespective of its distribution.

To begin with, let us observe that the problem of finding the equilibrium of the economy ultimately takes the form

$$\max_{\{s_i\geq 0\}} U\left(\boldsymbol{y}+\sum_{i=1}^N s_i\boldsymbol{\xi}_i\right) \tag{7}$$

This is self-evident from the consumer's viewpoint. But if (7) is solved, the firms' profit maximization is solved too. Indeed, taking the partial derivative of U with respect to s_i we find

$$\frac{\partial U}{\partial s_i} = \sum_{\mu=1}^P \frac{\partial u}{\partial x^\mu}\xi_i^\mu = \lambda\frac{\partial \pi_i}{\partial s_i} \tag{8}$$

where we used the fact that the consumer's utility maximization, subject to the budget constraint, implies that $\frac{\partial U}{\partial x^\mu} = \lambda p^\mu$ for some Lagrange multiplier $\lambda > 0$. This is the content of the first Welfare Theorem [5]: the market achieves an allocation of resources that is simultaneously optimal for all participants.

Eq. (7) is a standard problem in statistical mechanics, which can be solved by introducing a fictitious temperature, building a partition function with $-U$ in place of the Hamiltonian, and then letting the temperature to zero. In order to cope with the disorder (either hard- or soft-constrained), it is necessary to introduce replicas before averaging. This standard procedure generates macroscopic order parameters such as

$$q_{ab} = \frac{1}{N}\sum_{i=1}^N \Delta_i s_{ia}s_{ib} \tag{9}$$

where the indices $a, b = 1, \ldots, r$ run over replicas. The convexity properties of U ensure that the solution of the maximization problem is unique. This implies that the replicated partition function, which in the limit $N \to \infty$ can be computed via a simple saddle-point integration, is dominated by replica-symmetric saddles where

$$q_{a,b} = Q\delta_{a,b} + q(1-\delta_{a,b}). \tag{10}$$

(in spin-glass jargon, Q represents the self-overlap while q is the analog of the Edwards-Anderson order parameter). Here, we will leave routine calculations aside (details of a very similar calculation can be found in [10]). Let it suffice to say that after taking

the thermodynamic limit $N \to \infty$, the replica limit $r \to 0$, and the limit of zero temperature, one arrives at an expression of the form

$$\lim_{N\to\infty} \frac{1}{P} \left\langle \max_{\{s_i\geq 0\}} U\left(\boldsymbol{y} + \sum_{i=1}^{N} s_i \boldsymbol{\xi}_i\right)\right\rangle_{\boldsymbol{\xi}} = \mathrm{extr}_{\boldsymbol{\omega}}\ f(\boldsymbol{\omega}) \tag{11}$$

where $\boldsymbol{\omega}$ is a vector of order parameters and extr means that the solution is provided by the saddle point of the function f, i.e. by the vector $\boldsymbol{\omega}^*$ that solves the equation $\frac{\partial f}{\partial \boldsymbol{\omega}} = 0$. The function f depends in this case on the type of disorder $\boldsymbol{\xi}$ one is considering.

We shall first focus on the case of hard-constrained technologies, for which a well-behaved solution exists for all n, as long as $\eta > 0$. Such a solution exhibits two regimes, separated by a phase transition in the limit $\eta \to 0$. For small values of n, when few technologies are available, the economy possesses many unexploited productive opportunities and the introduction of new technologies increases the productivity of existing ones. For larger values of n, the economy enters a mature phase, where the productive sector is largely saturated by existing technologies.

In the soft-constrained case, instead, the economy becomes unstable as the average scale of production diverges for n larger than a critical η-dependent threshold. We will show explicitly that such an instability is related to the fact that, above this threshold, it is possible (with probability one) to combine existing technologies in order to produce outputs without using inputs.

3. Hard-constrained technologies

3.1. The effective-agent problems

If all technologies are strictly inefficient (i.e. if, for each i, $\sum_\mu \xi_i^\mu = -\eta\sqrt{\Delta_i}$), one finds that equilibria are described by the saddle point of

$$\begin{aligned} f(Q,\gamma,\chi,\widehat{\chi},\kappa,p) = {}& \frac{1}{2}nQ\widehat{\chi} - \frac{1}{2}\gamma\chi + \kappa p \\ & + \left\langle \max_{x\geq 0}\left[u(x) - \frac{1}{2\chi}\left(x - y + t\sqrt{nQ} + \kappa\right)^2\right]\right\rangle_{t,y} + \\ & + n\left\langle \max_{s\geq 0}\left[-\frac{1}{2}\Delta\widehat{\chi}s^2 + st\sqrt{\Delta(\gamma - p^2)} - s\eta p\sqrt{\Delta}\right]\right\rangle_{t,\Delta} \end{aligned} \tag{12}$$

where t is a unit Gaussian random variable, and averages over y and Δ are performed with distributions $\rho(y)$ and $g(\Delta)$, respectively. The representative-agent problem embodying the equilibrium structure of the multi-agent economy is rather clear. The first maximization problem on the r.h.s. can be interpreted as an effective-utility maximization by the consumer with respect to the consumption of a single "representative" commodity. The second one corresponds instead to effective-profit maximization by a "representative" company with respect to its operation scale. The two problems are interconnected in a non-trivial way by the other terms. The solutions $x^* \equiv x^*(t,y)$ and $s^* \equiv s^*(t,\Delta)$ of the effective problems as functions of the random variables (t,y) and (t,Δ), respectively, are given by

$$x^* \equiv x^*(t,y) \quad \text{such that} \quad \chi u'(x^*) = x^* - y + t\sqrt{nQ} + \kappa \tag{13}$$

and

$$s^*(t,\Delta) = \begin{cases} \dfrac{t\sigma - \eta p}{\widehat{\chi}\sqrt{\Delta}} & \text{for } t \geq \eta p/\sigma \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

Notice that the equilibrium consumption x^* is non-negative as long as $u'(x)$ diverges when $x \to 0$. Notice also that the dependence of the equilibrium scale s^* on Δ is described by

$$s^*(t,\Delta) = \frac{1}{\sqrt{\Delta}}\, s^*(t,1) \tag{15}$$

Moreover, by introducing the re-scaled variable $\widetilde{s} = s\sqrt{\Delta}$ the dependence of (12) on Δ disappears. This simple scaling will greatly simplify our discussion of the effects induced by heterogeneities in Δ (next section).

Using the shorthand $\sigma = \sqrt{\gamma - p^2}$ and a little algebra, the saddle-point conditions can be seen to take the following form:

$$p = \langle u'(x^*)\rangle_{t,y} \tag{16}$$

$$\sigma = \sqrt{\left[\langle u'(x^*)^2\rangle_{t,y} - \langle u'(x^*)\rangle_{t,y}^2\right]} \tag{17}$$

$$\widehat{\chi} = \frac{1}{\sqrt{nQ}}\,\langle t u'(x^*)\rangle_{t,y} \tag{18}$$

$$Q = \left\langle (s^*)^2 \Delta \right\rangle_{t,\Delta} \tag{19}$$

$$\chi = \frac{n}{\sigma}\left\langle t s^* \sqrt{\Delta}\right\rangle_{t,\Delta} \tag{20}$$

$$\kappa = p\chi + n\eta\left\langle s^*\sqrt{\Delta}\right\rangle_{t,\Delta} \tag{21}$$

In the next sections we shall discuss the generic properties of the above saddle-point problem and the corresponding phase structure.

3.2. Economic interpretation

The meaning of the macroscopic order parameters and of the saddle-point conditions can be easily found with at most some basic algebra.

- Equation (16) implies that the order parameter p is the average (relative) price, while σ (see (17)) measures (relative) price fluctuations. This is so because utility maximization under budget constraint gives $\frac{\partial u}{\partial x^\mu} = \lambda p^\mu$ with $\lambda > 0$ a Lagrange multiplier that can be set to 1 without any loss of generality. Evaluating this at the saddle point gives precisely (16).
- Averaging the expression in (13) over t and y and expressing κ via (21) one finds

  $$\langle x^*\rangle_{t,y} = \langle y\rangle_y - n\eta\left\langle s^*\sqrt{\Delta}\right\rangle_{t,\Delta} \tag{22}$$

  This equation can also be obtained by averaging the market clearing condition (4) over technologies taking (5) into account. Hence (22) can be seen as the macroeconomic analog of (4).

- It is possible to show via elementary manipulations [10] that

$$\langle u'(x^*)(x^* - y)\rangle_{t,y} = 0 \tag{23}$$

Remembering that $u'(x^*)$ is the relative price at equilibrium, this equation implies that the consumer saturates his/her budget when choosing his/her consumption. This law, known in economics as Walras' law [5], is a generic consequence of the microscopic setup of economic equilibria. The fact that it can be retrieved from our macroscopic saddle-point problem constitutes a significant consistency check.

3.3. Equilibrium distributions of operation scales, consumptions and (relative) prices

The distribution of operation scales $\mathcal{P}(s)$ provides important information, for it can be seen as a proxy for the distribution of firm sizes (either in terms of revenues or in terms of the number of employees). Its calculation in our model can be carried out with the help of (15). In fact, one has

$$\mathcal{P}(s) = \int_0^\infty g(\Delta) P(s|\Delta) d\Delta \tag{24}$$

where $P(s|\Delta)$ is the distribution of operation scales at fixed Δ. The latter can be derived from (14) and from the (Gaussian) distribution of t:

$$P(s|\Delta) = \langle \delta(s - s^*(t,\Delta))\rangle_t = (1-\phi)\delta(s) + \frac{\widehat{\chi}}{\sqrt{2\pi\sigma^2}}\, e^{-\frac{(\widehat{\chi} s\sqrt{\Delta}+\eta p)^2}{2\sigma^2}}\theta(s) \tag{25}$$

with

$$\phi = \left\langle \theta\left(t - \frac{\eta p}{\sigma}\right)\right\rangle_t = \frac{1}{2}\left[1 - \mathrm{erf}\left(\frac{\eta p}{\sigma\sqrt{2}}\right)\right] \tag{26}$$

the fraction of active firms (i.e. firms such that $s_i > 0$), which is independent of Δ. Now (15) implies that $P(s|\Delta) = \sqrt{\Delta} P(s\sqrt{\Delta}|1)$ so that

$$\mathcal{P}(s) = \frac{2}{s^3}\int_0^\infty k^2 g\left(\frac{k^2}{s^2}\right) P(k|1) dk \tag{27}$$

where we simply set $k = s\sqrt{\Delta}$. The behavior of $Q(s)$ for large s is related to the behavior of $g(\Delta)$ for small Δ. In particular, if $g(\Delta) \propto \Delta^\gamma$ for $\Delta \ll 1$, then $\mathcal{P}(s) \propto s^{-3-2\gamma}$ for $s \gg 1$, which agrees with the power law behavior found empirically for the distribution of firm sizes [11, 12].

The distribution of equilibrium consumptions at fixed endowments can be computed from (13) via the distribution of t. One gets

$$P(x|y) = \langle \delta(x - x^*(t,y))\rangle_t = \frac{1-\chi u''(x)}{\sqrt{2\pi n Q}}\, e^{-\frac{(x-y-\chi u'(x)+\kappa)^2}{2nQ}} \tag{28}$$

As before, the equilibrium distribution of consumptions is given by

$$\mathcal{P}(x) = \int_0^\infty \rho(y) P(x|y) dy \tag{29}$$

Likewise, it is also possible to compute, via the simple change of variable $p = u'(x)$, the equilibrium distribution of prices.

3.4. Limit $\rho(y) \to \delta(y-\overline{y})$

If the spread of initial endowments becomes vanishingly small, the amount of productive activity is expected, quite intuitively, to vanish too. In order to see this, one can proceed as follows. Let the initial endowment be expressed as

$$y = \overline{y} + \zeta \tag{30}$$

with ζ a (small) r.v. with zero average and $\overline{y}$ a positive number, and let

$$x^*(t,y) \equiv \overline{y} + \psi \qquad \psi \equiv \psi(t,\zeta) \tag{31}$$

Anticipating that $\psi \sim \zeta$ will be a small quantity when $\left\langle \zeta^2 \right\rangle_\zeta \ll 1$, we can write to leading order

$$u'(x^*) \simeq u'(\overline{y}) + \psi u''(\overline{y}). \tag{32}$$

Likewise,

$$p \simeq u'(\overline{y}) + u''(\overline{y}) \left\langle \psi \right\rangle_{t,\zeta} \tag{33}$$

$$\widehat{\chi} \simeq \frac{u''(\overline{y})}{\sqrt{nQ}} \left\langle t\psi \right\rangle_{t,\zeta} \tag{34}$$

$$\sigma^2 \simeq [u''(\overline{y})]^2 \left(\left\langle \psi^2 \right\rangle_{t,\zeta} - \left\langle \psi \right\rangle_{t,\zeta}^2 \right) \tag{35}$$

follow immediately from (16), (17) and (18). An expression for ψ can be obtained from (13) via (32):

$$\chi[u'(\overline{y}) + \psi u''(\overline{y})] \simeq \kappa + \psi - \zeta + t\sqrt{nQ} \tag{36}$$

Separating terms of order ψ^0 from those proportional to ψ, one has

$$\kappa = \chi u'(\overline{y}) \tag{37}$$

$$\psi(t,\zeta) = \frac{\zeta - t\sqrt{nQ}}{1 - \chi u''(\overline{y})} \tag{38}$$

Substituting (38) into (33), (34) and (35) and carrying out the averages one gets

$$p \simeq u'(\overline{y}) \tag{39}$$

$$\widehat{\chi} \simeq -\frac{u''(\overline{y})}{1 - \chi u''(\overline{y})} \tag{40}$$

$$\sigma^2 \simeq \widehat{\chi}^2 \left(\left\langle \zeta^2 \right\rangle_\zeta + nQ \right) \tag{41}$$

In order to calculate Q, we can use (41). Insertion into (14) yields $s^* = (t-\tau)s_0\theta(t-\tau)$, where

$$s_0 = \frac{\sigma}{\widehat{\chi}\sqrt{\Delta}} \simeq \sqrt{\frac{\langle \zeta^2 \rangle_\zeta + nQ}{\Delta}} \tag{42}$$

$$\tau = \frac{\eta p}{\sigma} \simeq \frac{\eta p}{\widehat{\chi}} \frac{1}{\sqrt{\langle \zeta^2 \rangle_\zeta + nQ}} \tag{43}$$

In turn, from (19) we obtain

$$Q \simeq \left(\left\langle \zeta^2 \right\rangle_\zeta + nQ \right) F(\tau) \quad \Rightarrow \quad Q \simeq \frac{\left\langle \zeta^2 \right\rangle_\zeta F(\tau)}{1 - nF(\tau)} \tag{44}$$

where

$$F(\tau) = \left\langle (t-\tau)^2\theta(t-\tau)\right\rangle_t \tag{45}$$

Now (44) and (45) together imply that $\tau \sim -\left\langle\zeta^2\right\rangle_\zeta^{-1/2}$ for small $\left\langle\zeta^2\right\rangle_\zeta$ so that, when $\left\langle\zeta^2\right\rangle_\zeta \to 0$, one can resort to an asymptotic expansion of $F(\tau)$ for $\tau \gg 1$ to calculate Q. This gives

$$Q \simeq \left\langle\zeta^2\right\rangle_\zeta^{3/2} \exp\left[-\frac{\eta^2[u'(\overline{y})]^2}{2[u''(\overline{y})]^2\left\langle\zeta^2\right\rangle_\zeta}\right] \tag{46}$$

So $Q \to 0$ for $\left\langle\zeta^2\right\rangle_\zeta \to 0$ with an essential singularity. In turn, by virtue of (41) and (26), respectively, $\sigma \to 0$ and $\phi \to 0$ for $\left\langle\zeta^2\right\rangle_\zeta \to 0$. Hence when the spread of initial endowments decreases, economic activity vanishes very rapidly and no market activity takes place when initial endowments are all equal.

3.5. Limit $\eta \to 0$: the phase transition

Before discussing the actual solutions of the saddle-point equations as a function of n and the corresponding economic picture, let us briefly investigate the limit $\eta \to 0$ of marginally efficient technologies, in which the critical properties of the model can be elucidated in detail. When $\eta = 0$, one sees from (26) that the fraction of active companies becomes $\phi = 1/2$, while from (14) one gets

$$s^*(t,\Delta) = \begin{cases} \dfrac{t\sigma}{\widehat{\chi}\sqrt{\Delta}} & \text{for } t \geq 0 \\ 0 & \text{for } t \leq 0 \end{cases} \tag{47}$$

The saddle-point conditions (19), (20) and (21) in turn take the simple form

$$Q = \frac{\sigma^2}{2\widehat{\chi}^2} \qquad \chi = \frac{n}{2\widehat{\chi}} \qquad \kappa = p\chi \tag{48}$$

One sees immediately that this provides a solution that must break down for $n > 2$, as there cannot be more than P active firms in the economy (i.e. ϕ must satisfy the condition $n\phi \leq 1$). Let us then focus on the limit $n \to 2^-$ by studying the case $0 < 2-n \ll 1$. When there is an almost complete set of technologies (i.e. when $n\phi \simeq 1$), goods can be transformed in almost all ways. Hence we expect that in equilibrium $x^* = \overline{y} + \psi$ with $\psi \equiv \psi(t)$ small. Then we can again use (32) and (13) to obtain an expression for ψ (as done in (36)). After simple algebra we find

$$\psi = \frac{2\widehat{\chi}(\overline{y}-y) + t\sigma\sqrt{2n}}{nu''(\overline{y}) - 2\widehat{\chi}} \tag{49}$$

Substituting this into (40) and (41) one obtains

$$\widehat{\chi} = -\frac{1}{2}u''(\overline{y})(2-n) \tag{50}$$

$$\sigma^2 = \frac{1}{2}[u''(\overline{y})]^2(2-n) \tag{51}$$

and thus, while the fluctuation of relative prices $\sigma \to 0$, because of (48), $Q \sim \frac{1}{2-n}$ and $\chi \sim \frac{1}{2-n}$. This means that the typical scale of production diverges as $\frac{1}{\sqrt{2-n}}$ when $n \to 2^-$.

A slightly more refined analysis is required to investigate the limit $\eta \to 0$ for $n > 2$. To simplify equations, for the remainder of this section we fix $\Delta = 1$ and resort to the notation introduced in the previous section, so that $s^* = (t-\tau)s_0\theta(t-\tau)$ with s_0 and τ as defined in the equalities in (42) and (43). Again anticipating that, because the set of technologies is almost complete and efficient, the equilibrium consumption levels $x^* = \overline{y} + \psi$ will fluctuate very little around the average $\overline{y}$ (i.e. $\psi \ll 1$), we can look for solutions of the saddle-point problem such that σ vanishes linearly with η (i.e. such that $\sigma \sim c\eta$) and with $s_0 = \sigma/\widehat{\chi}$ finite (i.e. such that $\widehat{\chi}$ vanishes linearly with η as well). In doing so, we will retain only the leading terms in η for the rest of this section. Now $\tau = p/c$, so that from (19) we have

$$Q = s_0^2 F(p/c) \tag{52}$$

with $F(\cdot)$ as in (45), whereas from (20) one gets

$$\chi = \frac{n s_0}{c\eta} H(p/c) \qquad H(x) = \frac{1}{2}\left[1 - \mathrm{erf}\left(\frac{x}{\sqrt{2}}\right)\right] \tag{53}$$

Expanding (13) as in (32), and using (52) and (53), one obtains

$$\psi \simeq \frac{c\eta}{n|u''(\overline{y})|H(p/c)}\left[\frac{1}{s_0}(y-\overline{y}) - t\sqrt{nF(p/c)}\right] \tag{54}$$

Inserting the above and (52) into (18) one finds

$$n^2 H(p/c)^2 \simeq 1 \tag{55}$$

This, via (26), confirms that $H(p/c) \equiv \phi \simeq 1/n$ (i.e. that technologies saturate the possible productions). Moreover, as $H(x) < 1/2$ for $x > 0$, it allows us to conclude that this solution holds for $n > 2$. Finally, it tells us that $c \to \infty$ for $n \to 2^+$ as $c \sim \frac{1}{n-2}$.

It remains to evaluate s_0. To this aim, we use (17). After minor algebraic manipulations, it yields

$$s_0 \sim \frac{1}{\sqrt{1 - nF(p/c)}} \tag{56}$$

For $n \to 2^+$ one has $c \to \infty$ and hence $F(p/c) \to 1/2$. So the typical scale of production (which is proportional to s_0) diverges as $\frac{1}{\sqrt{n-2}}$ for $n \to 2^+$.

In summary, we found that for $\eta \to 0$ a critical value $n_c = 2$ of n exists such that

$$\phi = \begin{cases} 1/2 & \text{for } n < n_c \\ 1/n & \text{for } n > n_c \end{cases} \tag{57}$$

while the typical scale of operations behaves as

$$\langle s^* \rangle \sim \frac{1}{\sqrt{|n - n_c|}} \qquad |n - n_c| \ll 1 \tag{58}$$

This shows that the behavior of the economy close $n = 2$ for $\eta \ll 1$ is totally analogous to that of a second-order phase transition in statistical physics.

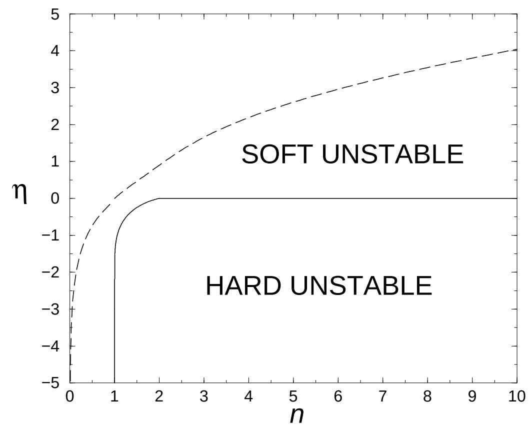

Figure 1. Phase diagram. Critical lines $n_c(\eta)$ for the economy with hard-constrained technologies (solid line) and soft-constrained technologies (dashed line). The economies are unstable for $n > n_c(\eta)$.

3.6. Transition to the Land of Cockaigne

The analysis performed so far can be easily extended to negative values of η. In this case, it is to be expected that the economy becomes unstable at certain critical points $n_c(\eta)$. The reason is that, for each fixed $\eta < 0$, combinations of technologies allowing to produce all goods without absorbing inputs (i.e. such that $\sum_i s_i \xi_i^\mu > 0$ for all μ) become possible with probability one when n is sufficiently large, larger than a critical η-dependent value $n_c(\eta)$. The equilibrium is stable for $n < n_c(\eta)$, whereas equilibrium operation scales diverge as $n \to n_c(\eta)$. Once this has happened, initial endowments become irrelevant, so that $n_c(\eta)$ is independent of $\rho(y)$. Rather, it only depends on the asymptotic behavior of the utility function $u(x)$. Fig. 1 shows the critical line $n_c(\eta)$ (solid line) separating the stable from the unstable phase for $u(x) = \log x$ and $g(\Delta) = \delta(\Delta - 1)$. This line was obtained by imposing the condition $n\phi = 1$ on the solution of the saddle point equations with $s_0 \to \infty$. As η decreases, $n_c(\eta)$ rapidly approaches 1, with an exponential behavior.

Notice that while this transition is continuous and characterized by divergences (i.e. $s_i \to \infty$) for $\eta \leq 0$ as $n \to n_c(\eta)$, it is discontinuous across the $\eta = 0$ line for $n > 2$. Indeed all scales of consumption $\langle x^* \rangle$ and production $\langle s^* \rangle$ remain finite as $\eta \to 0^+$ for $n > 2$ and diverge abruptly as soon as $\eta < 0$. This makes the phase diagram of this model very similar to that of Minority Games with market-impact correction [13]. Broadly speaking, both models deal with a system where a number N of agents compete for the exploitation of a fixed number P of resources and in both cases η tunes the efficiency of agent's behavior.

3.7. Solution in a typical example

Let us finally discuss the solutions of the saddle-point problem as a function of n. We have chosen, for simplicity, the following parameters:

$$g(\Delta) = \delta(\Delta - 1) \qquad u(x) = \log x \qquad \rho(y) = e^{-y} \tag{59}$$

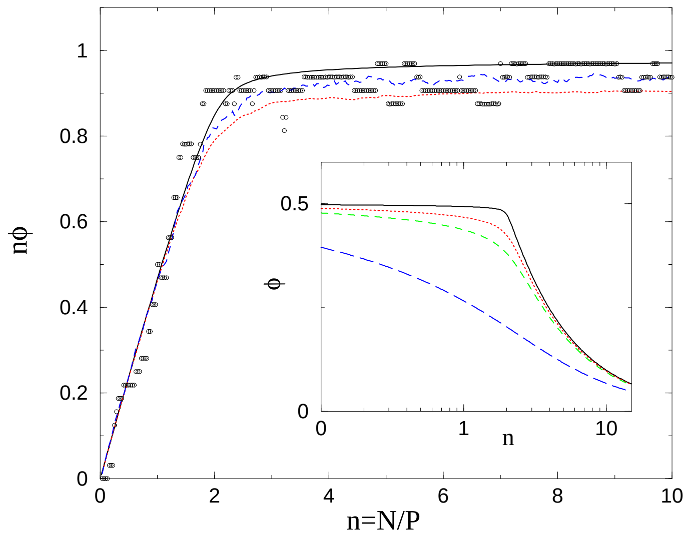

Figure 2. Economy with hard-constrained technologies. Behaviour of $n\phi$ ($\phi =$ fraction of active companies) at equilibrium as a function of n for $\eta = 0.05$ and $g(\Delta) = \delta(\Delta - 1)$: analytical prediction (continuous line), computer experiments with $P = 16$ (dotted line) and for $P = 32$ (dashed line) averaged over 100 disorder samples. Dots represent results of a single realization of the technologies. Inset: ϕ vs n for $\eta = 0.01, 0.05, 0.1, 0.5$ (top to bottom)

Correspondingly, we have obtained numerically the solution of Eq.s (16–21) as a function of n for different values of η. The resulting economic observables, representing equilibrium values, have been compared to those of equilibria computed numerically via a simple gradient descent algorithm for a toy economy with $P = 16$ and $P = 32$ commodities.

The fraction of active firms ϕ (or $n\phi$) is shown in Fig. 2. The two regimes described by (57) appear distinctly for low η. For $n < n_c = 2$, ϕ is roughly equal to 1/2. Here, about a half of the existing technologies are profitable. For $n > n_c$, instead, $\phi \simeq 1/n$, i.e. the number of active firms saturates the number of goods and the economy becomes extremely selective. This picture persists for larger values of η, although the transition appears to be more and more blurred as η increases.

In Fig. 3 we display the equilibrium behaviour of some economically relevant quantities. Let us discuss the behavior for $n < n_c = 2$ first. One sees that the average scale of production increases when n grows. This can be interpreted by saying that all firms benefit from the introduction of new technologies (i.e. from an increase of n) if the market is not too selective. In parallel, relative price fluctuations decrease as n increases, as does the average level of consumption, signalling that firms are managing the transformation of abundant goods into scarce ones. The distribution of consumptions preserves the character of the original distribution of endowments for low n [10], while relative fluctuations stay roughly unchanged.

When n is close to n_c, operation scales become larger and larger as η decreases (i.e. as technologies become more and more efficient) and ultimately develop a singularity at n_c in the marginally efficient limit $\eta \to 0$ (see (58)). The fluctuations of relative consumptions start to drop (the sharper the lower is η), as the distribution

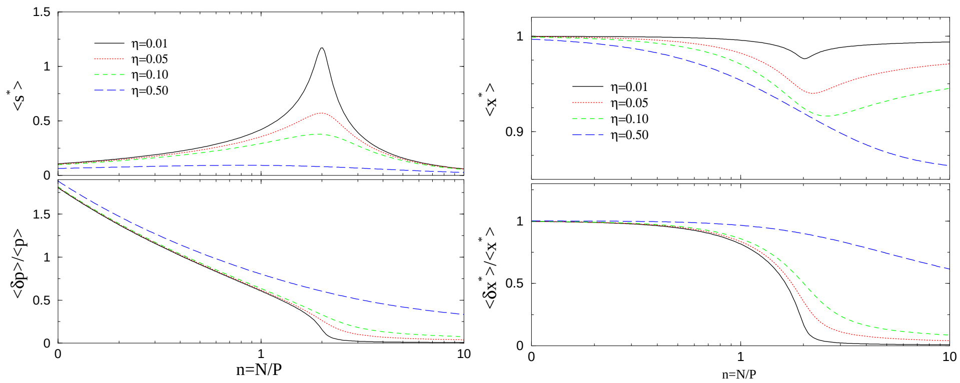

Figure 3. Economy with hard-constrained technologies. Behaviour of some relevant economic observables at equilibrium with $g(\Delta) = \delta(\Delta-1)$. Top left panel: average scale of operation. Bottom left panel: relative price fluctuations. Top right panel: average consumption level. Bottom right panel: relative fluctuations of the consumption level. Values of η are as indicated.

of consumptions becomes more and more peaked around the mean value. Identifying abundant (or scarce) goods becomes increasingly hard.

In the saturated regime $n > n_c$, technological innovations ($N \to N+1$) lead to a decrease in the average operation scale, i.e. new profitable technologies punish existing ones. Firms cannot take advantage of the spread between scarce and abundant goods any longer, and technologies are subject to a much stronger selection. On the other hand, the average consumption starts growing with n, as is expected in a competitive economy that selects highly efficient technologies.

It is interesting to notice (see [10] for a more detailed economic discussion) that while technological innovations (i.e. increases of N at fixed P) improve the global economic picture for $n < n_c$, when $n > n_c$ a global improvement requires broadening the range of disposable commodities (i.e. an increase of P at fixed N, or a decrease of n). Within the obvious limitations of this model, this seems to suggest that an economy driven by technological innovation and products creation may evolve 'spontaneously' toward the critical point n_c under its agents' maximizing pressures.

4. Soft-constrained technologies

4.1. The effective-agent problem and its typical solutions

Let us now turn our attention to the case of soft-constrained technologies. The analog of (12) in this case is

$$f(Q,\gamma,\chi,\widehat{\chi},m,\widehat{m}) = \frac{1}{2}nQ\widehat{\chi} - \frac{1}{2}\gamma\chi - nm\widehat{m} \tag{60}$$

$$+n\left\langle \max_{s\geq 0}\left[-\frac{1}{2}\Delta\widehat{\chi}s^2 + (\widehat{m}+t\sqrt{\gamma})\,s\sqrt{\Delta}\right]\right\rangle_{t,\Delta}$$

$$+\left\langle \max_{x\geq 0}\left[u(x) - \frac{1}{2\chi}\left(x-y+t\sqrt{nQ}+\eta nm\right)^2\right]\right\rangle_{t,y} \tag{61}$$

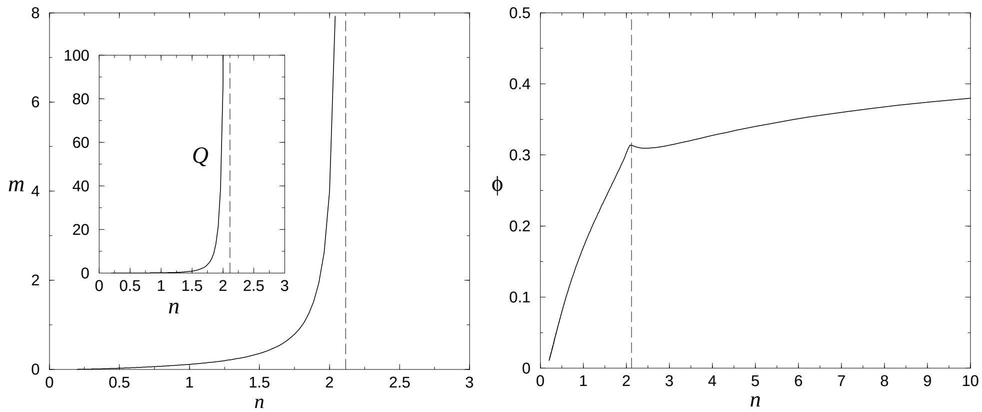

Figure 4. Economy with soft-constrained technologies. Solution of the saddle-point equations as a function of n for $\eta = 1$ and $g(\Delta) = \delta(\Delta - 1)$. Left panel: average scale of operation m and self-overlap Q. Right panel: fraction of active companies. The dashed vertical lines mark the position of the critical point n_c calculated analytically.

where we used the same notation for averages and variables as in Sec. 3. The effective problems for firms and consumers are easily identifiable and can be analyzed along the lines followed in Sec. 3.1. We jump straight to the saddle-point equations, which, using the same notation as in Sec. 3, now read

$$m = \left\langle s^* \sqrt{\Delta} \right\rangle_{t,\Delta} \tag{62}$$

$$Q = \left\langle (s^*)^2 \Delta \right\rangle_{t,\Delta} \tag{63}$$

$$\chi = \frac{n}{\sqrt{\gamma}} \left\langle t s^* \sqrt{\Delta} \right\rangle_{t,\Delta} \tag{64}$$

$$\gamma = \left\langle u'(x^*)^2 \right\rangle_{t,y} \tag{65}$$

$$\widehat{m} = -\eta \left\langle u'(x^*) \right\rangle_{t,y} \tag{66}$$

$$\widehat{\chi} = \frac{1}{\sqrt{nQ}} \left\langle t u'(x^*) \right\rangle_{t,y} \tag{67}$$

It is clear that m represents the average scale of operations while $\widehat{m}$ is connected to the average (relative) price and γ is related to price fluctuations. The above equations can again be solved numerically upon varying n. Adopting the same choices as in (59), one obtains the behaviour shown in Fig. 4. Strikingly different macroscopic properties emerge from this model. In particular, the average scale of operations diverges at a critical point that for $\eta = 1$ is roughly $n_c \simeq 2.1$. It is clear that such an economy is completely unrealistic. When m diverges, the average relative price drops to zero, so that arbitrarily large consumptions can be achieved (the reader can verify that relative price fluctuations are nevertheless finite above n_c). We will show in the next section that this behaviour is due to the fact that for $n > n_c$ technologies can be combined in such a way that $\sum_i s_i \xi_i^\mu > 0$ for all μ, i.e. that it is possible to produce every good without consuming.

4.2. The critical point

We will now proceed by calculating the typical volume of configuration space where technologies can be combined to yield $\sum_i s_i \xi_i^\mu > 0$ for all μ. For a fixed disorder realization $\boldsymbol{\xi}$, this volume is given by

$$V(\boldsymbol{\xi}) = \int \prod_{\mu=1}^{P} \theta\left(\sum_{i=1}^{N} s_i \xi_i^\mu\right) d\boldsymbol{s} \tag{68}$$

For $n < n_c$, it is to be expected that $V(\boldsymbol{\xi})$ is non-zero only with an exponentially small (in N) probability, whereas the probability of it being finite becomes of order 1 at n_c. The typical volume is obtained from the quenched average of $V(\boldsymbol{\xi})$:

$$V_{\text{typ}} = \exp \langle \log V(\boldsymbol{\xi}) \rangle_{\boldsymbol{\xi}} \tag{69}$$

We will focus in particular on the quantity

$$S = \lim_{N\to\infty} \frac{1}{N} \log V_{\text{typ}} \tag{70}$$

(analogous to a zero-temperature entropy in statistical physics). From the above discussion, we expect S to be negative for $n < n_c$ and positive for $n > n_c$, so that the critical point n_c can be determined from the condition $S = 0$.

In principle, n_c depends on η and $g(\Delta)$. Notice, however, that because of the statistics of the ξ_i^μ's, a re-scaling $s_i \to s_i\sqrt{\Delta_i}$ makes the integral independent of Δ_i's, apart from a trivial factor $\prod_i \frac{1}{\sqrt{\Delta_i}}$. Hence we shall set $\Delta_i = 1$ for all i without any loss of generality.

In order to ensure that integrals are well-defined, we introduce a spherical constraint of the form $\sum_i s_i^2 = N$ and concentrate on

$$V(\boldsymbol{\xi}) = \int e^{\mathrm{i}\sum_\mu \widehat{x}^\mu (x^\mu - \sum_i s_i \xi_i^\mu) - \frac{1}{2}\widehat{y}(\sum_i s_i^2 - N)} d\boldsymbol{s} d\boldsymbol{x} d\widehat{\boldsymbol{x}} d\widehat{\boldsymbol{y}} \tag{71}$$

As before, the evaluation of the $\boldsymbol{\xi}$-average requires replication techniques. The calculation is standard and we don't detail it here. In the limit $N \to \infty$ it ultimately turns out that S is given by the saddle-point value of

$$h(q, m, \widehat{m}, r, y) = \frac{1}{2}qr - m\widehat{m} + \frac{1}{2}y^2 - \frac{1}{n}\log 2 \tag{72}$$

$$+\frac{1}{2}\log\frac{\pi}{2(y+r)} + \frac{1}{n}\left\langle \log\left[1 - \operatorname{erf}\left(\frac{\eta n m + t\sqrt{nq}}{\sqrt{2n(1-q)}}\right)\right]\right\rangle_t \tag{73}$$

$$+\frac{\widehat{m}^2 + r}{2(y+r)} + \left\langle \log\left[1 + \operatorname{erf}\left(\frac{\widehat{m} + t\sqrt{r}}{\sqrt{2(y+r)}}\right)\right]\right\rangle_t \tag{74}$$

n_c can be found by solving simultaneously the saddle-point equations $\frac{\partial h}{\partial q} = \cdots = \frac{\partial h}{\partial y} = 0$ with $n = n_c$ together with the condition that h vanishes for $n = n_c$. This leads to the critical line $n_c(\eta)$ shown in Fig. 1 (dashed line) for both positive and negative values of η. The numerical value obtained for $\eta = 1$, namely $n_c = 2.11524\ldots$, is in convincing agreement with the results of the previous section. We remark that $n_c \to 1$ as $\eta \to 0$. Indeed, when $n \geq 1$ there are enough random vectors $\boldsymbol{\xi}_i$ to span the entire P-dimensional space. On the other hand, $n_c \sim \eta^2$ for $\eta \gg 1$.

5. Conclusions

The model we analyzed here is clearly unrealistic in many respects. However it shows that tools of statistical physics of disordered systems allow us to characterize the collective statistical properties of a complex economy, beyond what naïve intuition and simple probabilistic arguments would suggest. It is reasonably easy to anticipate a change of behavior for $n \approx 2$ on the basis of simple geometric arguments [10]. It is hard however to generalize these insights to derive, for example, the full picture of the phase transition for $\eta \to 0$ we have described, or the behavior of the economy in the two phases.

It would be desirable to extend this approach to the case with many consumers or to analyze the effects induced by imperfect competition (oligopolies, see e.g. [5]). A further extension to time dependent economies may provide interesting insights into economic growth theory. In particular, one can think about technological improvement as a dynamics of the ξ_i^μ's occurring on time scales much longer than those over which agents achieve equilibrium. This would require a theory just moderately more sophisticated than the one presented here.

Acknowledgments

We are indebted to D Fiaschi, KG Maler, F Vega-Redondo and MA Virasoro for important discussions. This work was partially supported by the EU EXYSTENCE network and by the EU Human Potential Programme under contract HPRN-CT-2002-00319, STIPCO.

References

[1] Mézard M, Parisi G and Virasoro MA. *Spin glass theory and beyond*. World Scientific (Singapore, 1987)
[2] Hertz J, Krogh A and Palmer RG. *Introduction to the theory of neural computation*. Addison-Wesley (Redwood City, CA, 1991)
[3] Martin OC, Monasson R and Zecchina R 2001 *Theor. Comp. Sci.* **265** 3
[4] http://www.unifr.ch/econophysics/minority/
[5] Mas-Colell A, Whinston MD and Green JR. *Microeconomic theory*. Oxford University Press (New York, NY, 1995)
[6] Debreu G. *Theory of value*. Yale University Press (New Haven, CT, 1959)
[7] Foellmer H 1974 *J. Math. Econ.* **1** 51
[8] Durlauf SN 1999 *Proc. Nat. Acad. Sci.* **96** 10582
[9] Lancaster K. *Mathematical economics*. Dover (New York, NY, 1987)
[10] De Martino A, Marsili M and Pérez Castillo I 2003 Preprint [cond-mat/0309533]
[11] Okuyama K, Takayasu M and Takayasu H 1999 *Physica A* **269** 125
[12] Axtell RL 2001 *Science* **293** 1818
[13] De Martino A and Marsili M 2001 *J. Phys. A: Math. Gen.* **34** 2525